MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963

Towards a Theory of
Impasse-driven Learning

Technical Report PCG-1

Kurt VanLehn

Departments of Psychology and Computer Science
Carnegie-Mellon University,
Schenley Park,
Pittsburgh, PA 15213 U.S.A.

DEPARTMENT
of
PSYCHOLOGY

Carnegie-Mellon University

Towards a Theory of
Impasse-driven Learning

Technical Report PCG-1

Kurt VanLehn

Departments of Psychology and Computer Science
Carnegie-Mellon University,
Schenley Park,
Pittsburgh,PA 15213 U.S.A.

February 12, 1987

**DTIC
ELECTE**
OCT 2 2 1987
D

Running Head: Impasse-driven learning

A186643

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION / AVAILABILITY OF REPORT
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE	Approved for public release; Distribution unlimited
4. PERFORMING ORGANIZATION REPORT NUMBER(S) PCG-1	5. MONITORING ORGANIZATION REPORT NUMBER(S) Same as Performing Organization

6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Personnel and Training Research Office of Naval Research (Code 1142 PT)
6c. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, PA 15213		7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy St. Arlington, VA 22217-5000
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Same as Monitoring Organization	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0349

8c. ADDRESS (City, State, and ZIP Code)	10. SOURCE OF FUNDING NUMBERS 442a558-02/12-15-86			
	PROGRAM ELEMENT NO N/A	PROJECT NO. N/A	TASK NO. N/A	WORK UNIT ACCESSION NO N/A

11. TITLE (Include Security Classification)

Towards a theory of impasse-driven learning.

12. PERSONAL AUTHOR(S)   Kurt VanLehn

13a. TYPE OF REPORT Technical	13b. TIME COVERED FROM 86Jun15 TO 87Oct31	14. DATE OF REPORT (Year, Month, Day) 87 Feb 12	15. PAGE COUNT 34

16. SUPPLEMENTARY NOTATION

17.	COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Learning, skill acquisition, failure-driven learning, impasse-driven learning, impasse

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

See reverse side.

| 20. DISTRIBUTION / AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION
22a. NAME OF RESPONSIBLE INDIVIDUAL Susan Chipman	22b. TELEPHONE (Include Area Code) (202) 696-4322	22c. OFFICE SYMBOL 1142 PT

Abstract

Earlier work has shown that students often reach impasses while trying to use a procedural skill that they are acquiring. The step that they believe should be executed next cannot be performed. If they are in a test-taking situation, where they may not receive help, they perform a specific, simple kind of problem solving, called repair. This report speculates about what happens when impasses occur during instructional situations where help is available. The conjecture is that the help that students receive -- either from the teacher, from examining the textbook, or from other information sources -- is reduced to a sequence of actions that will get the students past the particular impasse that is preventing them from completing the exercise problem. This action sequence is generalized to become a new subprocedure. The new subprocedure is inserted into the old procedure at the location where the impasse occurred. The proposed learning process is called impasse-driven learning.

Table of Contents

Towards a Theory of Impasse-driven learning

Kurt VanLehn

1. Introduction

Learning is widely viewed as a knowledge communication process coupled with knowledge compilation process (Anderson, 1985). The communication process interprets instruction thereby incorporating new information from the environment into the mental structures of the student. Knowledge compilation occurs with practice. It transforms the initial mental structures into a form that makes performance faster and more accurate. Moreover, the transformed mental structures are less likely to be forgotten. At one time, psychology concerned itself exclusively with the compilation process by using such simple stimuli (e.g., nonsense syllables) that the affects of the communication process could be ignored. The work presented here uses more complicated stimuli, the calculational procedures of ordinary arithmetic. For such stimuli, the effects of the knowledge communication process cannot be ignored. It will be shown later that certain types of miscommunication can cause students to have erroneous conceptions. The long-term objective of the research reported here is to develop of theory of the neglected half of learning, knowledge communication. Consequently, whenever the term "learning" appears below, it is intended to mean knowledge communication.

Earlier work has shown that students often reach "impasses" while trying to use a procedural skill that they are acquiring. An impasse occurs when the step that they believe should be executed next cannot be performed. If they are in a test-taking situation, where they may not receive help, they perform a specific, simple kind of problem solving, called "repair." This chapter speculates about happens when impasses occur during instructional situations, where help is available. The conjecture is that the help that the student receives -- either from the teacher, from examining the textbook, or from other information sources -- is reduced to the sequence of actions that will get the student past the impasse. This action sequence is generalized to become a new subprocedure. The new subprocedure is inserted into the old procedure at the location where the impasse occurred. The proposed learning process is called impasse-driven learning.

The research presented here began with the "buggy" studies of Brown and Burton (1978). Those studies found that students of certain procedural skills, such as ordinary multicolumn subtraction, had a

surprisingly large variety of bugs (i.e., small, local misconceptions that cause systematic errors). Early investigations into the origins of bugs yielded a theory of procedural problem solving, named Repair Theory (Brown & VanLehn, 1980). Among other accomplishments, Repair Theory predicted the occurrence of certain patterns of short-term instabilities in bugs. These instabilities were subsequently found (VanLehn, 1982). Recent research has investigated the relationship between the curriculum, the students' learning processes and the acquisition of bugs. A learning theory, named Step Theory, has been added to Repair Theory, yielding an integrated explanation for the acquisition of correct and buggy procedures (VanLehn, 1983a; VanLehn, 1983b). Step Theory describes learning at a large "grain size." Given a lesson and a representation of what the student knows prior to the lesson, Step Theory predicts what the student's knowledge state will be after the lesson.

Recently, attention has turned toward describing learning at a finer grain-size. The object of the current research is to describe the student's cognitive processing during a lesson. The research strategy is to replace Step Theory by augmenting Repair Theory, which already provides a fine-grained account of problem solving processes during diagnostic testing sessions, so that the new theory provides a fine-grained account of learning. If this strategy succeeds, the cognitive processes will account for both problem-solving behavior and lesson-learning behavior. Thus the new theory will provide a more integrated account of cognition as well as describing learning behavior at a finer grain size. In order to make it easier to contrast old and new theories, the new one will be dubbed RT2, and the old theory, which is a conglomerate of Repair Theory and Step Theory, will be referred to as RT1.

This chapter introduces RT2. It is, for the most part, speculation. Unlike RT1, RT2 has not been implemented as computer simulation, nor has its internal coherence and empirical accuracy been scrutinized with competitive argumentation (VanLehn, Brown & Greeno, 1984). Although the ideas behind RT2 are simple extensions of the principles of RT1, they seem capable of explaining much about human behavior. Moreover, they relate to current research in machine learning and language acquisition. A discussion of RT2, even in its current underdeveloped form, should be at least timely, and perhaps interesting as well.

The chapter begins with a discussion of the task domain and the kinds of behavior one finds students displaying. It then introduces the old theory, RT1. Readers who are familiar with RT1 from earlier publications may safely skip sections 2 and 3. The remainder of the chapter presents RT2 and discusses its relationship to other work in cognitive science.

2. Learning elementary mathematical skills

The goal of this research is to develop a rigorously supported theory of learning by taking advantage of AI's new modelling power. The long term research strategy is to begin by studying a particular kind of cognition, then if all goes well, to test the theory's generality on other kinds of cognition. The initial studies focused on how elementary school students learn ordinary, written arithmetic calculations. The main advantage of arithmetic procedures, from a methodological point of view, is that they are virtually meaningless to most students (Resnick, 1982). Most students treat arithmetic procedures as arbitrary formal manipulations (i.e., "symbol pushing"). Although this may frustrate teachers, it allows psychologists to study a complex skill without having to model a whole world's worth of background knowledge.

This section introduces the domain. First it describes the instruction that students receive, and then it describes the behavior they produce. The theory's main job is to explain what kinds of mental structures are engendered by that instruction and how those structures guide the production of the observed behavior.

2.1. Learning from lesson sequences of examples and exercises

In a typical American school, mathematical procedures are taught incrementally via a lesson sequence that extends over several years. In the case of subtraction, there are about ten lessons in the sequence that introduce new material. The lesson sequence introduces the procedure incrementally, one step per lesson, so to speak. For instance, the first lesson might show how to do subtraction of two-column problems. The second lesson demonstrates three-column problem solving. The third introduces borrowing, and so on. The ten lessons are spread over about three years, starting in the late second grade (i.e., at about age seven). These lessons are interleaved with lessons on other topics, as well as many lessons for reviewing and practicing the material introduced by the ten lessons. In the classroom, a typical lesson lasts an hour. Usually, the teacher solves some problems on the board with the class, then the students solve problems on their own. If they need help, they ask the teacher, or they refer to worked examples in the textbook. A textbook example consists of a sequence of captioned "snapshots" of a problem being solved (see figure 1). Textbooks have very little text explaining the procedure, perhaps because young children do not read well. Textbooks contain mostly examples and exercises.

Take a ten to make 10 ones.	Subtract the ones.	Subtract the tens.
2 15	2 15	2 15 ·
~~3~~ ~~5~~	~~3~~ ~~5~~	~~3~~ ~~5~~
· 1 9	· 1 9	· 1 9
	6	1 6

Figure 1: A typical textbook example.

2.2. Describing systematic errors with "bugs"

The observable output of the students' learning process is their performance while solving exercise problems. Error data are a traditional measure of such performance. There have been many empirical studies of the errors that students make in arithmetic (Buswell, 1926; Brueckner, 1930; Brownell, 1935; Roberts, 1968; Lankford, 1972; Cox, 1975; Ashlock, 1976). A common analytic notion is to separate systematic errors from slips (Norman, 1981). Systematic errors appear to stem from consistent application of a faulty method, algorithm or rule. Slips are unsystematic "careless" errors (e.g., facts errors, such as 7-3=5). Since slips occur in expert performance as well as student behavior, the common opinion is that they are due to inherent "noise" in the human information processor. Systematic errors, on the other hand, are taken as stemming from mistaken or missing knowledge, the product of incomplete or misguided learning. Only systematic errors are used in testing the present theory. See Siegler and Shrager (1984) for a developmental theory of addition slips.

Brown and Burton (1978) used the metaphor of bugs in computer programs in developing a precise, detailed formalism for describing systematic errors. A student's errors are accurately reproduced by taking a formal representation of a correct procedure and making one or more small perturbations to it, such as deleting a rule. The perturbations are called bugs. A systematic error is represented as a list of one or more bugs. Bugs describe systematic errors with unprecedented precision. If a student makes no slips, then his or her answers on a test exactly match the buggy algorithm's answers, digit for digit. Bug data are the main data for testing this theory.

6

Burton (1982) developed an automated data analysis program, called Debuggy. Using it, data from thousands of students learning subtraction were analyzed, and 76 different kinds of bugs were observed (VanLehn, 1982). Similar studies discovered 68 bugs in addition of fractions (Shaw et al., 1982), several dozen bugs in simple linear equation solving (Sleeman, 1984), and 57 bugs in addition and subtraction of signed numbers (Tatsuoka & Baillie, 1982).

It is important to stress that bugs are only a notation for systematic errors and not an explanation. The connotations of "bugs" in the computer programming sense do not necessarily apply. In particular, bugs in human procedures are not always stable. They may appear and disappear over short periods of time, often with no intervening instruction, and sometimes even in the middle of a testing session (VanLehn, 1982). Often, one bug is replaced by another, a phenomenon called bug migration.

Mysteries abound in the bug data. Why are there so many different bugs? What causes them? What causes them to migrate or disappear? Why do certain bugs migrate only into certain other bugs? Often a student has more than one bug at a time--why do certain bugs almost always occur together? Do co-occurring bugs have the same cause? Most importantly, how is the educational process involved in the development of bugs? One objective of the theory is to explain some of these bug mysteries.

Another objective is to explain how procedural skills are acquired from multi-year curricula. This objective seems to require longitudinal data, where each student in the study is tested several times during the multi-year period. Such data is notoriously difficult to acquire. Bug data are readily available and nearly as good. The bug data discussed here were obtained by testing students at all stages in the curriculum. Thus, the bug data are like between-subjects longitudinal data. Instead of testing the same student at several times at different stages of his or her learning, different students at different stages are tested just once. As will be seen later, such cross-sectional data can perform nearly as well as longitudinal data in testing a learning theory, and yet they are much easier to collect.

3. An introduction to the model: Explaining Always-Borrow-Left

Most of the mental structures and processes proposed by the theory can be introduced and illustrated by going through an explanation for a certain subtraction bug, called Always-Borrow-Left. Students with this bug always borrow from the leftmost column in the problem no matter which column originates the borrowing. Problem A below shows the correct placement of borrow's decrement. Problem B shows the

bug's placement.

$$
\begin{array}{llll}
 & \overset{5}{} & & \overset{2}{} & & \overset{5}{} \\
\text{A.} & 3\ 6^15 & \text{B.} & 3\ 6^15 & \text{C.} & 6^15 \\
 & \underline{-\ 1\ 0\ 9} & & \underline{-\ 1\ 0\ 9} & & \underline{-\ 1\ 9} \\
 & 2\ 5\ 6 & & 1\ 6\ 6 & & 4\ 6
\end{array}
$$

Always-Borrow-Left is moderately common. In a sample of 375 students with bugs, six students had this bug (VanLehn, 1982). It has been observed for years (Buswell, 1926, pg. 173, bad habit number s27). However, this theory is the first to offer an explanation for it.

The explanation begins with the hypothesis that students use induction (generalization of examples) in learning where to place the borrow's decrement. All the textbooks used by students in our sample introduce borrowing using only two-column problems, such as problem C above. Multi-column problems, such as A, are not used. Consequently, the student has insufficient information to induce an unambiguous description of where to place the borrow's decrement. The correct placement is in the left-adjacent column, as in A. However, two-column examples are also consistent with decrementing the left-most column, as in B.

The next hypothesis of the theory is that when a student is faced with such an ambiguity in how to describe a place, the student takes a conservative strategy and saves all the relevant descriptions. When inducing from two-column problems (e.g., C), the student describes the borrow-from column as "a column that is both left-adjacent to the current column and the left-most column in the problem."

Suppose that our student is given a diagnostic test at this point in the lesson sequence and that the test contains borrowing problems of all kinds. Suppose the student is faced with solving problem D, below.

$$
\begin{array}{llll}
\text{D.} & \begin{array}{c} 3\ 6\ 5 \\ \underline{-\ 1\ 0\ 9} \end{array} & \text{E.} & \begin{array}{c} 3\ 6^15 \\ \underline{-\ 1\ 0\ 9} \end{array}
\end{array}
$$

The student starts to borrow, gets as far as E, and is suddenly stuck. The student's description of where to borrow from is ambiguous because there is no column that is *both* left-adjacent and the left-most column. In the terminology of the theory, getting stuck while problem solving is called reaching an *impasse*.

It is hypothesized that whenever students reach an impasse on a test, they engage in *local problem solving*. Local problem solving is just like classical puzzle solving (Newell & Simon, 1972) in that there is

an initial state, a desired final state, and state-change operators. Here, the initial state is being stuck, and the desired final state is being unstuck. Unlike traditional problem solving, the state-change operators of local problem solving don't change the state of the exercise problem. Instead, they change the *state of the interpreter* that is executing the procedure. The operators do things like pop the stack of goals or relax the criterion for matching a description to the exercise problem. They do not do things like writing digits on the test paper. Because the local problem solver modifies the state of the procedure's interpretation, it is a kind of *meta-level* problem solving. The sequences of meta-level operators that succeed in getting students unstuck are called *repairs*. Note that what is being repaired is, roughly speaking, the impasse. Repairs do not change the procedure. To put it in terms of Newell's (1980) problem space hypothesis, the procedure works in one problem space, and local problem solving works in a second problem space that is "meta" to the base problem space. Returning to our stuck student, three common repairs to the impasse are illustrated below.

$$\text{F.} \quad \begin{matrix} 2 \\ 3\ 6^15 \\ \underline{-1\ 0\ 9} \end{matrix} \qquad \text{G.} \quad \begin{matrix} 5 \\ 3\ 6^15 \\ \underline{-1\ 0\ 9} \end{matrix} \qquad \text{H.} \quad \begin{matrix} 3\ 6^15 \\ \underline{-1\ 0\ 9} \\ 6 \end{matrix}$$

In F, the student has relaxed the description of which column to borrow from by ignoring the restriction that the column be left-adjacent to the current column. The remaining restriction, that the column be the left-most column in the problem, has the student decrement the hundreds column, as shown in F. This is one repair. It generates the bug Always-Borrow-Left. Another repair is shown in G. Here, the student has relaxed the borrow-from description by ignoring the left-most requirement. The decrement is placed in the left-adjacent column, yielding G. This repair generates a correct solution to the problem. In H, the student has chosen to skip the borrow-from entirely, and go on to the next step in the procedure. This repair generates a bug that is named Borrow-No-Decrement-Except-Last, because it only executes a borrow-from when it is unambiguous where to place the decrement, and that occurs only when the borrow originates in the last possible column for borrow. To sum up, three different repairs to the same impasse generate two different bugs and a correct version of subtraction.

It was mentioned earlier that students' bugs are not like bugs in computer programs because students' bugs are unstable. Students shift back and forth among bugs, a phenomenon called bug migration. The theory's explanation for bug migration is that the student has a stable underlying procedure, but that the procedure is incomplete in such a way that the student reaches impasses on some problems. The student can apply any repair she can think of. Sometimes she chooses one repair, and sometimes she

chooses others. The different repairs manifest themselves as different bugs. So bug migration comes from varying the choice of repairs to a stable, underlying impasse. In particular, the theory predicts that the three repairs just discussed ought to show up as a bug migration. In fact, they do. Figure 1 is a verbatim presentation of a diagnostic test showing the predicted bug migration.

Figure 2: Verbatim presentation of a test by subject 8 of class 17 showing three repairs to the same impasse. On problems D, E and G, one repair generates the bug Borrow-No-Decrement-Except-Last. (N.B. The subject does not always use scratch marks to indicate borrowing.) On problems H and I, another repair generates the correct borrow-from placement. On problems K, M, N, P, Q, R and S, a third repair generates the bug Always-Borrow-Left. There are slips on problems D, P, Q and S. On problem R, a second kind of impasse occurs. While processing the hundred's column, the subject tries to decrement the zero in the ten thousand's column. A repair to this impasse ultimately leads to the answer of 2 in the hundred's column.

This discussion of the bug Always-Borrow-Left has illustrated many of the assumptions of the theory. First, procedures are the result of generalization of examples, rather than, say memorization of verbal or written recipes. The main evidence for this assumption is that there are accidental, visual characteristics of the examples, viz. the placement of the decrement, that a non-example source of instruction, such as a verbal recipe, would not mention. The appearance of these visual characteristics in the acquired procedure is evidence that they were learned by induction (see VanLehn (1986) for a full defense of this idealization).

A second assumption is that learning occurs in the context of a lesson sequence, and that many bugs are caused by testing students who are in the middle of the lesson sequence on exercise types that they have not yet been taught how to solve. Perhaps such bugs should be welcomed as signs of a healthy learning process that may eventuate in a correct understanding of the procedure. Such a view of bugs is radically different from the traditional view, which considers bugs to be "bad habits" that need to be remediated. On the other hand, the bad-habit view may be appropriate for older students, some of whom have bugs long after the lesson sequence has been completed (VanLehn, 1982).

Another set of assumptions involves the notions of interpretation, impasses and repairs. A particularly important hypothesis is that repairs occur at the meta-level and change only the state of the interpretation. This hypothesis predicts the existence of bug migration. In fact, this prediction was made before any evidence of bug migration had been found (Brown & VanLehn, 1980). The surprising success of this forecast and the fact that it is an almost unavoidable consequence of the hypothesis provide strong support for the theory.

4. The stable bug problem

Although some students' behaviors can be characterized as bug migrations, other students appear to have the same bug throughout a test. When such students are tested again two days later, they often have the same bug (VanLehn, 1982). Some students even show the same bug when tested twice six months apart (VanLehn, 1982). Such data encourage the interpretation that some students have learned their bugs. That is, their bugs have become a part of the knowledge structure they use to encode their procedure. Such relatively permanent bugs are called "stable" in order to differentiate them from bugs the may exist only for a short time, then migrate/change into other bugs[1].

Stable bugs present a problem for RT1. Repairs do not modify the core procedure, but instead modify the state of the interpreter that is executing the core procedure. After a repair has been accomplished and the interpreter is running again, there is no trace of the effects of repair on the core procedure. Bug migrations are explained by assuming that the students apply different repairs at different occurrences of

---

[1]The proportion of students whose errors are due to stable bugs varies significantly with the grade level. In one study, 49% of the third graders had stable subtraction bugs, vs. 27% of the fourth graders and 13% of the fifth graders (VanLehn, 1982). The variation is due to the fact that more older children know the correct algorithm: 19% of the third graders were bug-free, vs. 39% of the fourth-graders and 60% of the fifth-graders.

the impasse. In order to explain a stable bug, one must assume that the student chooses to apply the same repair every single time the impasse occurs. Intuitively, this seems quite unlikely.

One way to explain stable bugs within the RT1 framework is to assume that the set of possible repairs is different for different individuals. Some students may only know about one repair, so they always choose that repair at an impasse. They will appear to have a stable bug. However, this hypothesis has difficulties. There are stable bugs that can only be generated by assuming that the students have two different impasses, and that the student repairs the first one with one repair, and the second one with a different repair. Students with such bugs must know at least two repairs, yet they consistently choose the same one at each choice point. Assuming that different students have different repairs does not help explain such multi-impasse stable bugs.

4.1. The Patch Hypothesis

As another potential explanation of stable bugs, one could augment RT1 by assuming that there is some memory of previous episodes of impasses and repairs. Stable bugs are generated by assuming that the student recognizes the impasse as one that has occurred before, and recalls the repair that was selected before and employed successfully. To perform such recall, the student must have some memory of the impasse and the repair. That is, the student's knowledge of the skill must consist of a set of impasse-repair pairs in addition to the core procedure. Such pairs are called patches (Brown & VanLehn, 1980). Thus, if the students have a stable bug, then they have a patch for that impasse. If they don't have a patch, then the impasse may cause bug migrations.

There are problems with the hypothesis that the student's knowledge consists of patches as well as the core procedure. First, it seems inelegant and unparsimonious. Patches are, essentially, condition-action rules. The condition is a description of particular interpreter states (i.e., a certain kind of impasse). A patch's action is some modification to make to the interpreter's state. The core procedure is also made up of condition-action rules. The only differences between patches and the core procedure's rules is that the rules' conditions can test the external problem states (i.e., the state of a partially completed subtraction problem) and the rules' actions can modify the external problem state. That is, the patches operate exclusively on the interpreter's state, while core procedure's rules operate on the external problem state as well. Nonetheless, there are more similarities than differences between patches and rules. It would be parsimonious to combine them.

The second problem with patches is that they must be somewhat abstract in order to function properly. In order for the patch to apply to multiple occurrences of an impasse, it must be a *description* of the interpreter's state. Thus, if a patch is acquired from, say, the first occurrence of an impasse, then the condition half of the patch must be abstract. It must not mention details of the interpreter state that are idiosyncratic to this particular occurrence, such as the values of digits in the problem. Similarly, the repairs must also be abstract descriptions of the modifications that were performed to the interpreter's state. Consequently, acquiring a patch is not simply a matter of storing a state and a state-change. Rather, patch acquisition requires non-trivial abstraction.

A third, more technical problem with patches is that they do not interface well with the pattern matching component of the interpreter. In order to represent descriptions of the external problem state, the procedure employs patterns. Such patterns are just like the usual ones found in, for instance the conditions of production rules. They consist of sets of relations whose arguments are variables or constants. In order to employ such patterns, the interpreter must have a pattern matcher. The matcher tries to fit the pattern to the representation of the external problem state. If the pattern matches, then the objects matched by the variables are often "read" and become a part of the interpreter's state. We saw an instance of this earlier, in the discussion of the bug Always-Borrow-Left. A pattern is used to represent the idea that the place to borrow from is (1) the left-most column in the problem, (2) a column that is adjacent to the column that is the current focus of attention, and (3) a column that is left of the current focus of attention. Speaking very approximately, the pattern for this concept employs three relations, one for each of the constraints listed above. It has two variables: one for the current focus of attention, and one for the column to be borrowed from. The following is an informal presentation of the pattern:

    (Is-leftmost-column New-focus) &
    (Is-adjacent-to New-focus Current-focus) &
    (Is-left-of New-focus Current-focus)

If the pattern matches, the object that is matched to the New-focus variable, namely a particular column in the problem, becomes the focus of attention for the borrow-from subgoal. The bug Always-Borrow-Left is generated when this pattern fails to match. Such mismatching occurs when borrowing originates in the units column of problems with more than two columns. In such problems, there is no column that meets all three constraints. The bug is generated when the second one is relaxed, allowing the pattern to match and picking out the left-most column of the problem as the focus of attention for borrowing-from. This causes the student to borrow from the left-most column, which is exactly what the bug Always-Borrow-Left does.

If the patch hypothesis is correct, then it should be possible to build a patch for Always-Borrow-Left. The impasse half of the patch can be quite simple. It can achieve the appropriate degree of abstraction by merely referring to the pattern. The description in the condition-half of the patch would read: "The pattern that fetches the borrow-from column does not match." However, there are problems implementing the repair-half of the patch. The following paragraphs present three possible implementations, all of which fail.

The repair could also be expressed in terms of the pattern. It needs to say something like "relax the second relation of the pattern." However, if this is taken literally, it means actually modifying the pattern by removing the second relation from it. Such modifications change the procedure itself. This makes it hard to explain bug migration--one would have to assume that the relaxation repair puts the deleted relations back, for instance.

A second possibility for the Always-Borrow-Left patch involves interrupting the pattern matching process. In order to accomplish the requisite relaxation, the repair would have to interrupt the pattern matcher right when it was about to apply the second constraint of the pattern, and somehow cause the pattern matcher to skip over that relation. Expressing this repair as a patch is difficult. It would require a precise specification, at the theoretical level, of a pattern matching algorithm, thus embroiling the theory in a layer of irrelevant detail.

A third option for the Always-Borrow-Left patch is to include a revised pattern that has all the relations except the second one. The interpretation of this description is for the local problem solver to perform pattern matching using this pattern, and substitute the results into the interpreter's state just as if the original pattern had been matched. This option works, usually. However, it has the flaw that on some occasions, the pattern stored in the patch does not match. This causes an impasse *inside* the local problem solver. That is, there can be an impasse while a person is trying to fix another impasse. The local problem solver is running "meta" to the interpreter, trying to repair the interpreter's impasse. We could assume that there is a meta-meta level, where another local problem solver runs, trying to repair the impasse that occurred inside the meta-level local problem solver. Such "towers" of meta-level interpreters have begun to appear in AI (Smith, 1982), but their properties are largely unexplored at this time. It is probably best to avoid postulating such multi-level architectures of students until they are better understood computationally.

To sum up: there are three methods, for representing the repair-half of the patch: (1) modifying the core procedure by deleting a relation from the pattern, (2) having the repair cause the pattern matcher to ignore the relation, and (3) storing a substitute pattern in the patch and matching it from inside the local problem solver. Because all these methods have defects, it seems that patches cannot represent the stability of bug that, like Always-Borrow-Left, depend on pattern relaxation in their repairs. This is just one problem with the patch hypothesis. The others, mentioned earlier, are its lack of parsimony,. since patches are quite similar to rules, and the fact that non-trivial abstraction is required for patches to be acquired.

4.2. Representing Stable bugs with mal-rules

The new version of the theory, RT2, takes the position that there are no patches. The student's knowledge of the skill consists only of a procedure. In order to represent stable bugs, the core procedure has "mal-rules[2]." Mal-rules are identical in format and function to the core procedure's regular rules. The difference is only that they cause the student to answer incorrectly, rather than correctly. Furthermore, RT2 assumes that mal-rules are acquired by the same learning mechanism as regular rules. Regular rules are acquired by induction of the teacher's examples. Mal-rules are acquired by induction from the "mal-examples" produced by local problem solving. If a student does not induce a mal-rule from the mal-example, pehaps because he did not attend to the mal-example, then bug migration may occur. Thus, stable bugs occur when mal-rules are induced, and unstable bugs occur otherwise.

Mal-rules are a much more parsimonious solution to the stable bug problem than patches. They are identical to rules, and they are acquired by the same mechanism as rules. Thus, mal-rules escape the first objections raised against patches.

Another objection was that patches couldn't represent stable pattern-relaxation bugs. This objection is also taken care of by the mal-rule hypothesis. To illustrate how, consider the bug Always-Borrow-Left again. In the normal course of events, students are first taught borrowing with two-column problems. Later they are taught how to solve three-column borrow problems. Recall that after the first lesson, the pattern is over constrained:

---

[2]Derek Sleeman coined the term "mal-rule" for his method of describing bugs in an objective, theoretically neutral fashion Although the mal-rules of RT2 are interpreted as lying at a deeper, more psychologically plausible level, the use of the term seems just as descriptive of how the rules function.

```
(Is-leftmost-column New-focus) &
(Is-adjacent-to New-focus Current-focus) &
(Is-left-of New-focus Current-focus)
```

The second borrowing lesson shows that when borrowing originates in the units column of three-column problems, it is the tens column that one borrows from. The learning mechanism utilizes such examples to eliminate the first relation from the pattern. That is, the learning mechanism does pattern relaxation.

If, on the other hand, mal-examples had been presented where the hundreds column was borrowed from, then pattern relaxation would delete the second relation. Such mal-examples can be generated when the learner is tested between the first and second lessons on borrowing. The over-constrained pattern will causes impasses, and the repair of those impasses generates the mal-examples. On this account, stable bugs like Always-Borrow-Left seem to be caused, ironically, by learning from one's mistakes.

5. Learning occurs at impasses

The introduction of this chapter promised a description of a fine-grained learning process. Although the responsibilities of the learning process have been increased, by including the generation of mal-rules as well as rules, the large-grained description of the learning process has not yet been refined. This section ventures a finer-grained description.

If learning occurs as a result of local problem solving, then the learning processes is likely to be interwoven with the local problem solving process. The main hypothesis is that inductive learning occurs at impasses. The "at" is used here in two senses. Learning occurs only when an impasse occurs. If there is no impasse, there is no learning[3]. The second sense of learning "at" impasses is more subtle. When an impasse occurs, the student is "at" some place in the procedure. That is, the interpreter for the procedure is reading some part of the control structure of the procedure. The hypothesis is that the control location of the impasse is the place where the newly learned piece of procedure will be inserted. That is, if the control structure is visualized as layed out spatially, say as a tree, then the hypothesis that learning occurs "at" impasses takes on a spatial interpretation: the spatial location of the impasse is the place where the new subprocedure will be attached to the existing procedure. So, the hypothesis has two

---

[3]Knowledge compilation may occur without impasses, but that is not the kind of learning that the theory describes. The theory aims to describe initial learning, or knowledge communication. Also, even if there is an impasse, learning may not occur. For instance, the student may not attend to the instruction at that time. Thus, the claim can be restated more precisely as if there is knowledge communication learning, then it occurs at impasses.

independent aspects: (1) learning occurs "at" impasses in the temporal sense, and (2) learning occurs "at" impasses in the control stucture sense.

First, let's examine the implications of the temporal aspect. According to RT1, the only activity that occurs in response to impasses is repair. The goal of repair is merely to get the interpreter past the impasse in any way possible. In particular, repair is not concerned with answering the problem correctly. Consequently, repairs rarely modify the interpretation is such a way that the problems are solved correctly. However, the impasse-driven learning hypothesis is intended to explain the acquisition of regular rules as well as mal-rules. To do this, the theory must be amended to allow other activities in reponse to impasses. Once the history of this research is reviewed, it will be easy to see what those activities should be.

The bug data that initiated the theory were collected in testing situations. The students were asked to answer problems without help from their teacher, friends or textbooks. If they got stuck, they would have to rely on their own knowledge to get unstuck. Thus, they repaired. However, students are not always in test taking situations when they solve problems. Often, they solve practice exercises in class or at home. In such situations, help is permitted. Indeed, students are encouraged to ask for help if they get stuck. So, the second kind of activity that may occur at impasses is receiving help.

Help seems to be the source of information that allows correct rules to be learned at impasses. For instance, suppose that a student gets stuck while doing seat work. He raises his hand. His teacher comes over. He asks, "I got to here and got stuck. What am I supposed to do next?" The teacher shows him what to do, saying, "You do this, and then this, then this." This short sequence of actions is just what the student needs. Not only does it get him around the impasse, but it is an example of a new subskill. The student may abstract the actions, leaving behind details that are specific to the particular problem that is being solved, such as the numerical values of the digits. The abstracted actions become a new subprocedure, which the student can attach to his existing procedure. Thus, correct rules can be acquired by the same mechanism that acquires mal-rules, provided that there is some way of obtaining help at impasses.

The actual method of delivering help is probably of secondary importance. The student can obtain help by comparing his solution to his friend's solution. The comparison isolates a subsequence of actions illustrating the new subprocedure just as effectively as asking the teacher, but it may cost the student

some effort to make the comparison. With slightly more effort, the student could generate a subsequence by drawing a "near" analogy to a worked problem in the textbook (Anderson, Farrell, & Saurers, 1984). The student may even be able to generate the subsequence from a "far" analogy, given a little coaching from the teacher (Resnick & Omanson, 1987). For instance, some teachers might have the student think of the problem's base-ten numerals as piles of pennies, dimes and dollars. Under the teacher's prodding, the student maps the impasse over to the monetary representation, solves the impasse there without violating principles of fair exchange (e.g. a while ... change a dime for ten pennies), then maps those monetary manipulations back into paper-and-pencil actions in the written representation. Since the impasse was solved correctly in the monetary representation, the written analog of that solution should be a correct subsequence of actions. The point is that there are a variety of ways that a written action subsequence may be obtained. Demonstration, comparison, near analogy, and far analogy are only a few of the many possible ways, although they may be the most common ones.

As soon as one discovers that there are several kinds of inputs to a cognitive process, one wonders whether those differences make any difference. Does it matter whether the student receives help via individualized demonstrations, vs comparison etc. The simplest hypothesis is that it is the subsequence of actions that determines the contents of the subprocedure, and not the source for that subsequence of actions. If two methods of obtaining help yield the same subsequence of actions for the impasse, then the subprocedure that the student induces should be the same.

Like most simple hypotheses, this one is likely to be only half right, at best. The various methods of obtaining help may require different cognitive resources, and that may affect the inductive learning process. For instance, suppose far analogy takes longer and requires more problem solving of the student than attending to the teacher's demonstration. The heavier demands of far analogy could interfer with the retention of the interpreter's state at the time of the impasse. This may decrease or perturb inductive learning, because retaining (or reconstructing) the interpreter's state is necessary for determining where to attach the new subprocedure. This interference could be considered, however, a second order or "performance" factor.

The second aspect of the hypothesis that learning occurs "at" impasses concerns the place where the new subprocedure will be attached to the existing procedure. The hypothesis is that the attachment point will be the same as the location of the interpreter at the occurrence of the impasse. The notion of location is complicated by the fact that the student's procedures have a hierarchical control structure. That is, the

18

procedure has goals, which call subgoals, which call sub-subgoals and so on. There may even be recursion: a goal calling itself. Consequently, at the time of an impasse, there may be a whole stack of goals pending. However, it is the lowest goal, the one that is a subgoal of all the others, that is suffering from the impasse.[4] This has implications for where new subprocedures will be attached, given the hypothesis that subprocedures are attached at the place in the goal structure where the impasse occurred. Roughly speaking, new subprocedures will tend to be attached low in the goal tree. This prediction is a necessary implication of the hypothesis.

It is also a true prediction for the data available now. In order to predict the locations of subprocedure attachments as inferred from the arithmetic bug data, RT1 had an ad hoc hypothesis, called the lowest-parent hypothesis (VanLehn, 1983a). It simply stipulated that new subprocedures be attached as low as possible in the goal hierarchy. This hypothesis is no longer needed. The attachment points are predicted from an independently motivated hypothesis, viz., that learning occurs at impasses.

6. Implications for remediation

One of the most obvious facts about arithmetic is that remediation of bugs tends to work. Many students have arithmetic bugs when tested in the early grades (e.g., 49% in the third grade). The proportion decreases with grade level. The proportion of adults with bugs is much smaller. Apparently, students' bugs are being remediated somewhere in the educational process. The question addressed in this section is how this remediation takes place. Section 4.2 discussed how mal-rules are learned, and section 5 discussed how regular rules are learned. This section speculates about how regular, correct rules may be learned when they have to compete with mal-rules that were learned earlier.

Suppose a student enters a remedial session with a stable bug. The stability of the bug indicates that a mal-rule was learned during some prior episodes of local problem solving. Those impasses no longer occur because the mal-rule circumvents them. Consequently, the student can work all the relevant problems without reaching an impasse. As mentioned earlier, if there are no impasses, there is no

---

[4] Here is an informal proof that only the lowest goal can be stuck: Suppose that some goal other than the lowest goal is stuck. This means that there is some subgoal of it that is pending but not stuck. But if that lower goal is not stuck, then it can continue until it succeeds or fails. In either case it would be removed from the stack of pending goals. So the only goals left in the stack when interpretation is forced to stop are the stuck goal and all the goals that depend on its completion in order for them to complete, i.e., the supergoals of the stuck goal. In short, it is always the lowest goal on the stack that is the current locus of control when an impasse occurs.

learning. Yet, remediation works. How can the hypothesis be reconciled with the facts?

The simplest reconciliation is to assume that remediation occurs when the teacher stops the student just as the student makes a wrong move in the problem solving process. For instance, the teacher may be carefully watching the student do a borrow, and interrupts just as the student has placed a scratch mark where no scratch mark should occur. Suppose further that the student interprets the teacher's interruption in a similar fashion to an impasse. The student observes the subsequence of actions that the teacher suggests, abstracts them, and plugs them into the existing procedure as a new subprocedure. Treating interruptions in tutorial situations as impasses could extend the theory to cover remediation. Such interruptions are rare events, which explains why stable bugs are often found. However, if we postulate that such remediation is effective when it does occur, and that tutorial efforts persist for years, then the assumption also explains why stable bugs eventually disappear.

However, there is a small problem. Typically, the teacher's interruption occurs just after the first incorrect action instead of before it. Yet, that action is the result of a mal-rule which is running in place of the impasse. We want the new subprocedure to be placed where the impasse used to occur, just as if the instruction had been delivered then instead of now. This would cause the new subprocedure to be attached correctly. However, the teacher's interruption is too late. The impasse-place has been passed. In principle, the student could be asked to reason backwards in order to locate the impasse-place. However, a better remediation technique may be to have the student solve the same problem again (or a very similar one) and interrupt just before the incorrect action. This interruption would at the impasse place, or at least much closer. This is, of course, a testible suggestion concerning the effectiveness of two remediation strategies.

There is a more significant problem with the kind of remediation proposed so far. Suppose the interruption is completed, and the student has installed the new subprocedure at the impasse place. There is already another subprocedure attached there, the mal-rule. It was generated in response to that impasse. So both the rule (subprocedure) and mal-rule are attached at the same impasse place. Subsequently, when the student comes to that place in solving a problem, how will the student know whether to execute the rule or the mal-rule? Since both were constructed in order to handle the same impasse, both will be applicable. All other things being equal, the student will pick the rule half the time and the mal-rule the rest of the time. This predicts a new kind of instability, where buggy problem solving alternates with correct problem solving.

There is annecdotal evidence that this prediction might be close to the mark. The evidence concerns the spontaneous reappearance of bugs after their supposed remediation. Teachers have noted that when students with bugs are shown the correct procedure in a remedial session, they pick it up easily. They can solve dozens of problems successfully in the session. Apparently, they have learned the correct procedure. However, when tested again several weeks later, they are either back to using their old buggy procedure, or they are alternating between their old buggy procedure and the correct procedure, It is common to heard anecdotes about this phenomenon. Resnick and Omanson (1987) have carefully documented several cases of such bug regression in a study designed to investigate new remediation strategies. Bug regression occurred despite the fact that the remediation was particularly thorough. Nonetheless, Resnick and Omanson report that 60% of the students reverted to using their buggy procedures for answering written subtraction problems when tested about four weeks later.

Bug regression makes intuitive sense, given the cognitive process sketched above. Suppose the student has enough context during the remediation session to differentiate the newly learned rule from the mal-rule, which was learned some time ago. The similarity between the learning context and the application context allows the student to reliably differentiate the correct rule from the older rule, and thereby apply the new rule throughout the remediation session. However, at a later testing session, the context during the session may not be similar enough to the remediation session that the student can recall which rule is the one to use. This would cause the student to be uncertain which rule was correct; they might alternate rules in order to maximize their test score. Another possibility is that the mal-rule was learned during a testing session. The present context, another testing session, may be more similar to the context in which the mal-rule was learned than the context in which the correct rule was learned. This may cause the student use the mal-rule exclusively. Thus, depending on when the mal-rule was learned, the students may either apply the buggy procedure exclusively, or they may alternate between the correct procedure and their buggy one. The predictions of impasse-driven learning are in accord with the phenomenon of bug regression.

7. General Discussion

Many cognitive theories of learning have hypothesized that learning was some kind of automatic phenomenon. Mental activity leaves a trace that somehow makes it easier to perform that activity the next time. Automatic learning has been the dominate paradigm for the last few decades of psychology, if not longer. Particular examples of this kind of learning for skill acquisition are Anderson's (1983) ACT*

theory and Anzai and Simon's (1979) theory of learning by doing. These theories feature automatic learning of new material (i.e., task-specific productions) by repeated usage of older material (i.e., weaker, more general productions).

Automatic learning theories have begun to draw fire from computer scientists who have noted that the lack of control over what is learned causes the system to acquire vast quantities of useless knowledge (Minton, 1985). For instance, Roger Schank (1982) has rejected automatic learning as a totally impossible way to acquire common sense knowledge about, e.g., how to dine in a restaurant. He points out that most mundane thinking is so banal and disconnected that to remember it all would be pointless and a poor model of our introspective experience of learning. To put it in a phrase, automatic learning would generate mental clutter.

Impasse-directed learning does not generate mental clutter. Learning only occurs when the current knowledge base is insufficient. Moreover, it is not just any incompleteness that causes learning. The incompleteness must be relevant enough to the person's affairs that it actually generates an impasse. The person's problem solving must require a piece of knowledge that isn't there. Consequently, one learns only when there is a need to learn. Mental clutter is avoided, and only pertinent knowledge is acquired.

7.1. Related models of skill acquisition

Impasse-directed learning is a species of failure-driven learning. Failure-driven learning is a common theoretical idea in the learning literature. For instance, in Wexler and Culicover's (1980) theory of language acquisition, whenever the learning model can't understand a sentence, it randomly deletes a rule from its grammar, or it makes a change in an existing rule by randomly choosing from a small class of legal perturbations. Their theory is typical of a class of learning theories where negative reinforcement of an internal kind causes a more-or-less random change in the learner's knowledge. Impasse-driven learning is more specific than these theories in that it postulates exactly what kinds of negative reinforcement cause learning (i.e., impasses) and exactly what kinds of changes the learner makes to its knowledge. The impasse-driven learning hypothesis is a new member of the class of failure-driven learning theories.

The idea of impasse-driven learning is central to the SOAR architecture (Laird, Rosenbloom, & Newell, 1986). SOAR is a production system. When SOAR reaches an impasse, it does some problem solving at

the meta-level. As it returns from that problem solving, it automatically builds a new production rule whose conditions are exactly the conditions pertaining at the impasse and whose actions are the results of the meta-level problem solving that was just completed. If ever those conditions occur again, the production will fire, thus saving SOAR the effort of reaching an impasse and resolving it at the meta-level. SOAR's authors call this kind of learning "chunking" and the productions built this way are called chunks.

Soar's authors claim that chunking is the only kind of learning that people do. However, this claim is not very restrictive, because the SOAR architecture allows arbitrary meta-level problem solving at impasses. The chunking mechanism saves the results, but the programmer can generate those results any way she wants by writing the appropriate problem solving into SOAR's meta-level. RT2 will be more specific than that. The theory will describe in detail the meta-level problem solving *qua* learning that occurs at impasses.

The impasse-driven learning hypothesis has appeared in the literature on formal theories of natural language acquisition. Robert Berwick (1985) has developed a theory of how English syntax is learned. His theory is strikingly similar to RT2, despite the fact that the two theories were developed independently. Berwick assumes that a person has a grammar and a parser. The grammar and parser are analogous, respectively, to the procedure and interpreter postulated by RT2. As internal state, Berwick's parser employs a stack and some other temporary structures. These have analogs in RT2 as well. In Berwick's theory, the parser can get stuck because no grammar rules apply (the analog of reaching an impasse in RT2). One of four actions is taken. All four actions modify only the parser's internal state just like repairs would. Two of Berwick's four "repairs" have nearly exact analogs to the repairs found in RT2. So the architecture postulated by Berwick for understanding English is nearly isomorphic to the one we have arrived at for following procedures.

Berwick goes on to state his version of the impasse-driven learning hypothesis: grammar rules are induced when the parser gets stuck. Which rules are induced depends on external information, namely, a perceptually given understanding of the sentence. To put it intuitively, if the child can't understand a sentence, she figures out what it meant from context, then invents a rule that would both get her parser unstuck and be consistent with the sentence's meaning. This process is analogous to the one we postulate, except that the typical learner may appeal to a blackboard, a Dienes Blocks algorithm (given an implementation of Resnick and Omanson's suggestion) or some other source of information about the skill, rather than inferring its meaning from context.

7.2. Summary

Impasse-driven learning has been put forward as a conjecture about how students learn procedural skills. It employs the meta-level architecture proposed by Repair Theory. It postulates additional processes that run at the meta-level. When an impasse occurs, the student can either repair or seek help; both processes run at the meta-level and fix the problem of being at the impasse. When the impasse is fixed, the student can choose either to abstract the actions taken to resolve it, or not. In general, inducing a new subprocedure from the actions taken at the impasse will result in either a correct subprocedure, if help was sought, or a buggy subprocedure, if repair was used.

Impasse-driven learning seems to make the right sort of predictions about bugs and their stability. It predicts that bugs can migrate as well as be stable over long periods. It predicts that remediation of bugs will appear effective at the end of the remediation session, but that bugs will tend to reappear over time.

Impasse-driven learning also seems to correctly predict the shape of cognitive structures that are built by learning. It predicts that new subprocedures will be attached as deeply as possible in the goal-subgoal hierarchy of the student's procedures.

Impasse-driven learning is a form of failure-driven learning. Failure-driven learning has traditionally been advanced as more cognitively economical than automatic learning, its traditional opponent hypothesis, in that it predicts that new knowledge is acquired only when there is a need for that knowledge. Automatic learning tends to generate mental clutter--cognitive structures of little or no relevance to subsequent thinking.

Impasse-driven learning seems to have great potential generality. It has been investigated in a powerful general learning system, SOAR (Laird, Rosenbloom, & Newell, 1986). It has been shown capable of learning English grammar (Berwick, 1985). The future of this hypothesis seems quite bright indeed.

24

References

Anderson, J. R. (1983). *The Architecture of Cognition.* Cambridge, MA: Harvard.

Anderson, J. R. (1985). *Cognitive psychology and its implications.* New York: Freedman.

Anderson, J. R., Farrell, R., & Saurers, R. (1984). Learning to program in LISP. *Cognitive Science, 8,* 87-129.

Anzai, Y. & Simon, H.A. (1979). The theory of learning by doing. *Psychological Review, 86,* 124-140.

Ashlock, R. B. (1976). *Error Patterns in Computation.* Columbus, OH: Bell and Howell.

Berwick, R. (1985). *The Acquisition of Syntactic Knowledge.* Cambridge, MA: MIT Press.

Brown, J. S., & Burton, R. B. (1978). Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science, 2,* 155-192.

Brown, J. S. & VanLehn, K. (1980). Repair Theory: A generative theory of bugs in procedural skills. *Cognitive Science, 4,* 379-426.

Brownell, W. A. (1935). Psychological considerations in the learning and teaching of arithmetic. In W. D. Reeve (Ed.), *The Teaching of Arithmetic.* New York, NY: Teachers College, Bureau of Publications.

Brueckner, L. J. (1930). *Diagnostic and remedial teaching in arithmetic.* Philadelphia, PA: Winston.

Burton, R. B. (1982). Diagnosing bugs in a simple procedural skill. In D. H. Sleeman & J. S. Brown (Eds.), *Intelligent Tutoring Systems.* New York: Academic. 157-183.

Buswell, G. T. (1926). *Diagnostic studies in arithmetic.* Chicago, IL: University of Chicago Press.

Cox, L. S. (1975). Diagnosing and remediating systematic errors in addition and subtraction computation. *The Arithmetic Teacher, 22,* 151-157.

Laird, J. E., Rosenbloom, P. S., and Newell, A. (1986). Chunking in Soar: The anatomy of a general learning mechanism. *Machine Learning, 1*(1), 11-46.

Lankford, F. G. (1972). *Some computational strategies of seventh grade pupils.* Charlottesville, VA: University of Virginia.

Minton, S. (1985). Selectively generalizing plans for problem-solving. In *Proceedings of IJCAI 85.* Los Altos, CA: IJCAI,

Newell, A. & Simon, H. A. (1972). *Human Problem Solving.* Englewood Cliffs, NJ: Prentice-Hall.

Norman, D. A. (1981). Categorization of action slips. *Psychological Review, 88,* 1-15.

Resnick, L. (1982). Syntax and semantics in learning to subtract. In T. Carpeter, J. Moser & T. Romberg (Ed.), *A cognitive perspective.* Hillsdale, NJ: Erlbaum.

Resnick, L. B., & Omanson, S. F. (1987). Learning to understand arithmetic. In R. Glaser (Ed.), *Advances in instructional psychology*. Hillsdale, NJ: Erlbaum.

Roberts, G. H. (1968). The failure strategies of third grade arithmetic pupils. *The Arithmetic Teacher, 15*, 442-446.

Schank, R. (1982). *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge: Cambridge University Press.

Shaw, D. J., Standiford, S. N., Klein, M. F. & Tatsuoka, K. K. (1982). *Error analysis of fraction arithmetic-- selected case studies* (Tech. Report 82-2-NIE). University of Illinois. Computer-based Education Research Laboratory.

Siegler, R. S. & Shrager, J. (in press). Strategy choices in addition: How do children know what to do? In C. Sophian (Ed.), *Origins of Cognitive Skill.* .

Sleeman, D. H. (1984). Basic algebra revised: A study with 14-year olds. *International Journal of Man-Machine Studies,* .

Sleeman, D. (1984b). An attempt to understand students' understanding of basic algebra. *Cognitive Science, 8*, 387-412.

Smith, B.C. (1982). *Reflection and semantics in a procedural language* (Technical Report MIT-TR-272). M.I.T. Laboratory for Computer Science.

Tatsuoka, K. K. & Baillie, R. (1982). *Rule space, the product space of two score components in signed- number subtraction: an approach to dealing with inconsistent use of erroneous rules* (Tech. Report 82-3-ONR). University of Illinois, Computer-based Education Research Laboratory. Urbana, IL.

VanLehn, K. (1982). Bugs are not enough: Empirical studies of bugs, impasses and repairs in procedural skills. *The Journal of Mathematical Behavior, 3*(2), 3-71.

VanLehn, K. (1983). *Felicity conditions for human skill acquisition: Validating an AI-based theory* (Tech. Report CIS-21). Xerox Palo Alto Research Center.

VanLehn, K. (1983). Human skill acquisition: Theory, model and psychological validation. In *Proceedings of AAAI-83*. Los Altos, CA: Morgan Kaufman,

VanLehn, K. (1983). Arithmetic procedures are induced from examples. In J. Hiebert (Ed.), *Conceptual and Procedural Knowledge: The Case of Mathematics*. Hillsdale, NJ: Erlbaum.

VanLehn, K., Brown, J. S., & Greeno, J. G. (1984). Competitive argumentation in computational theories of cognition. In W. Kinsch, J. Miller & P. Polson (Ed.), *Methods and Tactics in Cognitive Science*. Hillsdale, NJ: Erlbaum.

26

Wexler, K. & Culicover, P. (1980). *Formal principles of language acquisition.* Cambridge, MA: MIT Press.

Dr. Phillip L. Ackerman
University of Minnesota
Department of Psychology
Minneapolis, MN 55455

Dr. Beth Adelson
Department of Computer Science
Tufts University
Medford, MA 02155

Dr. Robert Ahlers
Code N711
Human Factors Laboratory
Naval Training Systems Center
Orlando, FL 32813

Dr. Ed Aiken
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Robert Allen
Temple University
School of Business Administration
Department of Computer and
Information Sciences
Philadelphia, PA 19122

Dr. James Algina
University of Florida
Gainesville, FL 32605

Dr. John Allen
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. William E. Alley
AFHRL/MOT
Brooks AFB, TX 78235

Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Thomas H. Anderson
Center for the Study of Reading
174 Children's Research Center
51 Gerty Drive
Champaign, IL 61820

Dr. Steve Andriole
George Mason University
School of Information
Technology & Engineering
4400 University Drive
Fairfax, VA 22030

Technical Director, ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Alan Baddeley
Medical Research Council
Applied Psychology Unit
15 Chaucer Road
Cambridge CB2 2EF
ENGLAND

Dr. Patricia Baggett
University of Colorado
Department of Psychology
Box 345
Boulder, CO 80309

Dr. Eva L. Baker
UCLA Center for the Study
of Evaluation
145 Moore Hall
University of California
Los Angeles, CA 90024

Dr. Meryl S. Baker
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Isaac Bejar
Educational Testing Service
Princeton, NJ 08450

Leo Beltracchi
United States Nuclear
Regulatory Commission
Washington DC 20555

Dr. Mark H. Bickhard
University of Texas
EDB 504 ED Psych
Austin, TX 78712

Dr. John Black
Teachers College
Columbia University
525 West 121st Street
New York, NY 10027

Dr. Arthur S. Blaiwes
Code N711
Naval Training Systems Center
Orlando, FL 32813

Dr. Robert Blanchard
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. R. Darrell Bock
University of Chicago
NORC
6030 South Ellis
Chicago, IL 60637

Dr. Jeff Bonar
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Richard Braby
NTSC Code 10
Orlando FL 32751

Dr. Jomills H. Braddock II
Center for the Social
Organization of Schools
The Johns Hopkins University
3505 North Charles Street
Baltimore, MD 21218

Dr. Robert Breaux
Code N-095R
Naval Training Systems Center
Orlando, FL 32813

Dr. Ann Brown
Center for the Study of Reading
University of Illinois
51 Gerty Drive
Champaign, IL 61280

Commanding Officer
CAPT Lorin W. Brown
NROTC Unit
Illinois Institute of Technology
3300 S Federal Street
Chicago, IL 60616-3793

Dr. John S. Brown
XEROX Palo Alto Research
Center
3333 Coyote Road
Palo Alto, CA 94304

Dr. John Bruer
The James S. McDonnell
Foundation
University Club Tower, Suite 1610
1034 South Brentwood Blvd
St Louis, MD 63117

Dr. Bruce Buchanan
Computer Science Department
Stanford University
Stanford, CA 94305

Dr. Patricia A. Butler
OERI
555 New Jersey Ave, NW
Washington, DC 20208

Dr. Tom Cafferty
Dept of Psychology
University of South Carolina
Columbia, SC 29208

Dr. Joseph C. Campione
Center for the Study of Reading
University of Illinois
51 Gerty Drive
Champaign, IL 61820

Joanne Capper
Center for Research into Practice
1718 Connecticut Ave, NW
Washington, DC 20009

Dr. Susan Carey
Harvard Graduate School of
Education
337 Gutman Library
Appian Way
Cambridge, MA 02138

Dr Pat Carpenter
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr John M Carroll
IBM Watson Research Center
User Interface Institute
P O Box 218
Yorktown Heights, NY 10598

LCDR Robert Carter
Office of the Chief
of Naval Operations
OP-01B
Pentagon
Washington, DC 20350-2000

Dr Alphonse Chapanis
8415 Bellona Lane
Suite 210
Ruxton Towers
Baltimore, MD 21204

Dr Davida Charney
English Department
Penn State University
University Park, PA 16802

Dr Paul R Chatelier
OUSDRE
Pentagon
Washington, DC 20350-2000

Dr Michelene Chi
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Dr L J Chmura
Computer Science and Systems
Code 7590
Information Technology Division
Naval Research Laboratory
Washington, DC 20375

Mr Raymond E Christal
AFHRL/MOE
Brooks AFB, TX 78235

Dr William Clancey
Stanford University
Knowledge Systems Laboratory
701 Welch Road, Bldg C
Palo Alto, CA 94304

Dr Charles Clifton
Tobin Hall
Department of Psychology
University of
Massachusetts
Amherst, MA 01003

Dr Allan M Collins
Bolt Beranek & Newman, Inc
50 Moulton Street
Cambridge, MA 02138

Dr Stanley Collyer
Office of Naval Technology
Code 222
800 N Quincy Street
Arlington, VA 22217-5000

Dr Lynn A Cooper
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

LT Judy Crookshanks
Chief of Naval Operations
OP-112G5
Washington DC 20370-2300

Phil Cunniff
Commanding Officer, Code 7522
Naval Underwater Warfare Engineering
Keyport, WA 98345

Dr Cary Czichon
Intelligent Instructional ...ems
Texas Instruments AI Lab
P O Box 660245
Dallas, TX 75266

Brian Dallman
3400 TTW/TTGXS
Lowry AFB CO 80230-5000

Dr Natalie Dehn
Department of Computer and
Information Science
University of Oregon
Eugene, OR 97403

Dr Gerald F DeJong
Artificial Intelligence Group
Coordinated Science Laboratory
University of Illinois
Urbana, IL 61801

Geory Delacote
Directeur de L informatique
Scientifique et Technique
CNRS
15 Quai Anatole France
75700 Paris FRANCE

Dr Thomas E DeZern
Project Engineer, AI
General Dynamics
PO Box 748
Fort Worth, TX 76101

Dr Andrea di Sessa
University of California
School of Education
Tolman Hall
Berkeley, CA 94720

Dr R K Dismukes
Associate Director for Life Sciences
AFOSR
Bolling AFB
Washington, DC 20332

Dr Stephanie Doan
Code 6021
Naval Air Development Center
Warminster, PA 18974-5000

Defense Technical
Information Center
Cameron Station Bldg 5
Alexandria VA 22314
Attn TC
(12 Copies)

Dr Thomas M Duffy
Communications Design Center
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr Richard Duran
University of California
Santa Barbara, CA 93106

Edward E Edwards
CNATRA N301
Naval Air Station
Corpus Christi TX 78419

Dr John Ellis
Navy Personnel R&D Center
San Diego, CA 92252

Dr Jeffrey Elman
University of California
San Diego
Department of Linguistics C-008
La Jolla, CA 92093

Dr Susan Embretson
University of Kansas
Psychology Department
426 Fraser
Lawrence KS 66045

Dr Randy Engle
Department of Psychology
University of South Carolina
Columbia SC 29208

Dr William Epstein
University of Wisconsin
W J Brogden Psychology Bldg
1202 W Johnson Street
Madison WI 53706

ERIC Facility-Acquisitions
4833 Rugby Avenue
Bethesda MD 20014

Dr K Anders Ericsson
University of Colorado
Department of Psychology
Boulder CO 80309

Dr Steven Zornetzer
Office of Naval Research
Code 114
800 N Quincy St
Arlington, VA 22217-5000

END
DATE
FILMED

JAN

1988